

METHOD OF TRANSFERRING DATA BY TRANSMITTING LOWER ORDER AND UPPER ORDER MEMORY ADDRESS BITS IN SEPARATE WORDS WITH RESPECTIVE OP CODES AND START INFORMATION

This is a continuation of application Ser. No. 08/667,293, filed Jun. 19, 1996, now abandoned, which is a continuation of application Ser. No. 08/484,917, filed Jun. 7, 1995 now abandoned, which is a divisional of application Ser. No. 08/381,015, filed Jan. 30, 1995 now abandoned, which is a continuation of application Ser. No. 07/848,421, filed Mar. 6, 1992 now abandoned.

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to the transmission of data between devices coupled to a high speed bus system. More particularly, the present invention relates to the packet format transmitted across a high speed bus system and the processing of the same by devices coupled to the bus.

2. Art Background

A computer bus is utilized for communication of information among master and slave devices coupled to the bus. Generally, a bus comprises a plurality of transmission lines to which the devices are coupled. Address, control, and data information are multiplexed over the transmission lines forming the bus. The information is communicated across the bus in many different formats. One such format is a packet format in which data is bundled in packets for transmission on the bus across multiple clock cycles. An example of a bus which utilizes packets is described in PCT international patent application number PCT/US91/02590 filed Apr. 16, 1991, published Oct. 31, 1991, and entitled *Integrated Circuit I/O Using a High Performance Bus Interface*.

An example of a packet issued by a requesting device is illustrated in FIG. 1. Using bus lines BusCtl and BusData [7:0], in the first bus cycle the type of bus access and the master device (i.e., requesting device) is provided. In the second through sixth bus cycles the address of the requested data and the block size are provided.

However, as the speed of transmission of information on the bus increases, the speed required of the receiving devices to process the packet needs to also increase in order to reduce the latency and realize the advantages of the increased speed of transmission across the bus. Furthermore, it is desirable to decrease the die space consumed while maintaining full functionality at the bus interface.

SUMMARY AND OBJECTS OF THE INVENTION

It is therefore an object of the present invention to provide a packet format which enables the receiving device to decrease the latency when the packet is processed.

It is an object of the present invention to provide a packet format which enables a receiving device to initiate access operations as quickly as possible based upon the address provided in the packet.

It is an object of the present invention to increase the speed for determining packet collisions on the bus and notifying devices of the occurrence of the same.

It is further an object of the present invention to provide a packet format for transmission across a high speed bus in which the block size decoding at the receiving device is

09559835-042600

simplified thereby increasing the speed at which the receiving device processes the information.

It is an object of the present invention to provide a packet format for transmission on a high speed bus which enables the die space consumed on the device receivers to be reduced.

A high speed bus system in which at least one master device, such as a processor, and at least one DRAM slave device are coupled to the bus. An innovative packet format and device interface which utilizes a plurality of time and space saving features in order to decrease the die size of the device receiver and decrease the overall latency on the bus is provided. In the preferred embodiment the request packet is transmitted on ten multiplexed transmission lines, identified as BusCtl and BusData [8:0]. The packet is transmitted over six sequential bus cycles, wherein during each bus cycle, a different portion of the packet is transmitted. The lower order address bits are moved ahead of the higher order address bits of the memory request. This enables the receiving device to process the memory request faster as the locality of the memory reference with respect to previous references can be immediately determined and page mode accesses on the DRAM can be initiated as quickly as possible. The type of memory access is arranged over a plurality of clock cycles; placing the more critical bits first. The count of blocks of data requested is arranged to minimize the number of bit positions in the packet used and therefore the number of transmission lines of the bus and the number of bus receiver contacts on the receiving device. This helps minimize the amount of die space required on the chip to process the block count information. The number of blocks is encoded in a manner to decrease the die space consumed in the receiver as well as to simplify the decoding by the receiver, thereby increasing the speed along critical paths and decreasing latency during the processing of the request packet.

BRIEF DESCRIPTION OF THE DRAWINGS

The objects features and advantages of the present invention will become apparent to one skilled in the art when reading the following detailed description in which:

FIG. 1 illustrates a prior art packet format utilized in a high speed bus.

FIG. 2 is a block diagram illustration of an illustrative high speed bus structure.

FIG. 3 illustrates a preferred embodiment of the packet format of the present invention.

FIG. 4 illustrates another embodiment of the packet format of the present invention in which active collision detection of packets is performed.

FIG. 5a and FIG. 5b illustrate the decrease in length of the carry chain by organization of the information in the packet format.

FIGS. 6a and 6b illustrate the innovative encoding of bits for generation of byte masks utilized.

FIGS. 7a, 7b, 7c and 7d illustrate the innovative encoding technique employed for byte transfers of varying lengths.

DETAILED DESCRIPTION

The request packet format is designed for use on a high speed multiplexed bus for communication between master devices such as processors and slave devices, such as memories and, in particular, dynamic random access memories (DRAMs). The bus carries substantially all address, data and control information needed by the master devices for

005240-528560

communication with the slave devices coupled to the bus. The bus architecture includes the following signal transmission lines: BusCtl, BusData [8:0], as well as clock signal lines and power and ground lines. These lines are connected in parallel to each device as illustrated in FIG. 2.

The processors communicate with the DRAMs to read and write data to the memory. The processors form request packets which are communicated to the DRAMs by transmitting the bits on predetermined transmission lines at a predetermined time sequence (i.e., at predetermined clock cycles). The bus interface of the DRAM receiver processes the information received to determine the type of memory request, the address of the memory request and the number of bytes of the transaction. The DRAMs then perform the memory operation indicated by the request packet.

The memory address consists of the row address which is used during the row address strobe (RAS) in the DRAM and the column address which is used during the column address strobe (CAS) in the DRAM. The DRAMs have the capability to operate in normal RAS access mode or in page mode. When operable in page mode, if a subsequent request to access data is directed to the same row, the DRAM does not need to wait for receipt of the row address and to assert RAS, as RAS has been asserted during the previous memory access. Thus, the access time for this data is shortened. For further discussion regarding page mode DRAMs, see Steve L. Gumm, Carl T. Dreher, *Unraveling the Intricacies of Dynamic RAM*, Electronic Design News, pp. 155-165 (Mar. 30, 1989).

The request packet format further helps to improve the performance of the DRAMs in response to memory requests for page mode access. The DRAMs use the lower order portion of the memory address as the column address bits. This provides a locality of reference such that bytes of memory which are logically contiguous will be physically contiguous in the memory space. The resultant effect is that a greater number of logically contiguous bytes of memory are also physically contiguous and the frequency of page mode accesses is increased.

To further increase the access speed for a memory request, the lower order bits are placed at the beginning of the packet. This is illustrated in FIG. 3, where address bits Address [9:2] are placed in the first word of the packet and bits Address [17:8] are placed in the second word of the packet. By placing the lower order bits at the beginning of the packet, those memory accesses performed in page mode can be processed at least two cycles earlier further increasing the performance of the memory accesses.

As the lower order bits of the memory address are placed in the first two words of the packet, little room is left at the beginning of the packet for op code bits, op[3:0], which identify the type of memory operation to be performed (e.g., page mode access). However, as the memory operation type needs to be determined in order to perform the memory operation, the op code bits need to be transmitted early in the packet. In the packet format of the present invention, the BusCtl line and the most significant bit of the Data signal line, BusData[8], are utilized to transmit the op code bits. The bits are transmitted within the first 4 words of the packet, coincident with the transmission of the memory address. Preferably the memory operation types are coded in such a manner that the bits transmitted coincident with the lower order bits of the memory address indicate whether a page mode memory operation is to be performed.

At bus cycle zero, the BusCtl line is used to indicate the start of the packet.

When multiple devices are transmitting on a bus, the possibility of packet collisions exists. Many different techniques are employed to avoid the concurrent transmission of multiple packets on a bus. For example, the master devices keep track of all pending transactions, so that each master device knows when it can send a request packet and access the corresponding response. However, the master devices will occasionally transmit independent request packets during the same bus cycle. Those multiple requests will collide as each master device drives the bus simultaneously with different information, resulting in scrambled request information. Prior art techniques for detecting and responding to collision detection generally have been found to be too slow for high speed buses. Thus a mechanism for the detection of packet collisions on high speed buses is needed.

Typically two types of collisions will occur: those which are completely aligned in which two or more master devices start transmission at exactly the same cycle, and those which are unaligned in which two or more master devices start transmission at different cycles which are close enough together to cause overlap of the request packets. In PCT international patent application number PCT/US91/02590 filed Apr. 16, 1991, published Oct. 31, 1991, and entitled *Integrated Circuit I/O Using a High Performance Bus Interface*, collisions were detected by the master devices and signals indicating the existence of the collision were subsequently sent by the master devices to the slave devices. This technique requires the master devices to process the detection of a collision and drive the bus to notify the slave devices in a very short period of time. To eliminate need for the master device to notify the slave device of the collision, the master devices and the slave devices detect and process the existence of a collision in parallel.

Additional bits of the packet are preallocated to store a code which identifies the master device transmitting the packet. This is illustrated in the packet format shown in FIG. 4. At bus cycles 4 and 5, the processor device code, Master[3:0] is transmitted. If two master devices issue packet requests starting at the same bus cycle, the master device code, Master[3:0], will be logically ORed together resulting in a different code. This is detected in parallel by the master devices and slave devices which are monitoring the bus signal lines. The slave devices immediately respond by discarding the packets received and an arbitration is performed to determine priority of master device access to the bus for retransmission of the request packets.

An unaligned collision condition arises when a first master device issues a request packet at cycle 0 and a second master device issues a later packet starting, for example, at cycle 2 of the first request packet, thereby overlapping the first request packet. This will occur as the bus operates at high speeds, and the logic in the second master device may not be fast enough to detect a request initiated by the first master at cycle 0 and delay its own request. As the collision occurs during the later clock cycles of the first packet, it is critical that the slave device receiving the request know of the collision before completion of transmission of the request packet so that the packet can be discarded before the slave device responds to the request. The high speed of the bus increases the difficulty of the master device timely notifying the slave device of the occurrence of a collision. Therefore a second innovative collision detection mechanism is used for unaligned collisions. The BusCtl signal line is used at the first bus cycle to indicate the start of a packet. Referring to FIG. 4, BusCtl is now also utilized at predetermined bus cycles for collision detection which increases the speed at which a collision is detected and responded to.

The BusCtl line is monitored by the slave devices as well as the master devices to detect collisions. The BusCtl line at bus cycles at which a subsequent packet may be initiated are normally driven to a low or off state. In the present embodiment, packets are initiated on even clock cycles; therefore the BusCtl line during clock cycles 2 and 4 are normally driven to a low or off state. When a collision occurs, the BusCtl line at one or both cycles will be driven to an on or one state due to the overlap of the data, specifically the start packet signal of a subsequent packet. Both master devices and slave devices monitor BusCtl for information such as the start of the packet. Upon detecting an on or one state at cycles 2 and/or 4, the slave devices immediately know that a collision has occurred and eliminate the packet being received. Thus there is no requirement for the master device to notify the slave device, no delay in responding to a collision and no possibility that the transmission of the packet is completed before the slave device is notified of the collision.

The master devices also monitor the BusCtl signal line for the occurrence of a packet collision. Upon detection of an on state at cycle 2 and/or 4, the transmitting master devices will arbitrate access to the bus and retransmit the packets to ensure accurate transmission of the packets. Thus, the technique described enables the slave devices to immediately detect the occurrence of a collision and discard the packets before the slave devices respond to the requests.

The encoding and decoding of the number of bytes or "count" for a memory operation also plays a significant role in decreasing the latency of processing the transaction. In the high speed bus which utilizes the packet format of the present invention, a balance is achieved between the number of bits required to encode the byte count for the memory transaction and the complexity of logic at the receiver interface of the memory device and the speed of operating the same. Referring to FIG. 3, a total of eight bits are used, Count[7:0]. Although the bits could have been transmitted during the same cycle across parallel transmission lines, the bits have been deliberately organized across two sequential bus cycles and transmitted across adjacent transmission lines. By placing the information on adjacent transmission lines in sequential bus cycles, the amount of wiring required to move the data received in the receiver from the bus input to the receiver logic which determines the count is decreased as there is simply shorter distances between the data inputs. This is illustrated by the block diagrams set forth in FIGS. 5a and 5b.

FIG. 5a is a simplified representation of a physical implementation of a slave device bus interface. In this illustration, count bits [7:2] are transmitted across the bus during one clock cycle on parallel bus lines. The bits are received at the inputs of the bus interface 100, 105, 110, 115, 120 and 125. Once these bits are received, the bits are processed through logic components (not shown) which provide a counter function which counts the number of quadbytes to be transferred. The implementation of this counter requires a carry chain to be built. The decrease in length of the carry chain by placing the count bits on adjacent transmission lines in sequential bus cycles is conceptually illustrated by FIGS. 5a and 5b. The length of the wire 130 needed to form the carry chain for the single clock cycle transmission as shown in FIG. 5a is much greater than the length of wire 135 used to form the carry chain for the bits transmitted sequentially and in parallel as shown in FIG. 5b. The decrease in wire length minimizes the amount of die area required at the receiver and further affects the speed of data along critical paths and thus the latency for decoding the count information.

To simplify the implementation of the receivers of the memory devices as well as reduce the die size of the receiver and decrease the latency for processing bus transactions, data is accessed in the memory in groups of four bytes, referred to herein as "quadbytes". Although the discussion below is directed to the transmission of data in quadbytes, it will be obvious to one skilled in the art from reading the following discussion that the concepts can be extended to any multiple byte organization.

The count bits not only identify the number of bytes to be transmitted starting at the identified memory address, but also the location of the bytes in the quadbyte transmitted. For example, the memory address of the request identifies a location within a quadbyte. To eliminate those bytes not requested, the memory device will mask out the unwanted bytes. The mask is also determined from the count value. In the preferred embodiment, the memory device masks out unwanted bytes during write transactions. During read transactions all bytes of the quadbyte are transferred across the bus. The processor then eliminates those bytes of the first and last quadbyte received which were not requested. This is preferred because this simplifies the implementation of the data path inside the memory devices. For example, in the preferred embodiment, this eliminates the need for a space consuming and time consuming data alignment network to insure proper sequencing of individual bytes. The additional logic that would be required to support the masking and other functions, such as the data alignment network, at the memory devices contributes to increasing the complexity of the chip as well as increasing the die size. However, it should be realized that the memory device can be configured to perform masking operations on both read and write transactions in order to eliminate any unwanted bytes of quadbytes prior to transmission across the bus.

A processor wishing to formulate a memory request will have an internal byte address, `MasterAddress[35:0]` and an internal byte length, `MasterCount[7:0]` for the data to be transferred pursuant to the request. Using offset-by-one encoding, the convention used is as follows: `MasterCount[7:0]=00000000` indicates one byte and `MasterCount[7:0]=11111111` indicates 256 bytes. The processor converts these internal values into the values for the request packet according to the following:

`Address[35:0]=MasterAddress[35:0]`

`Overflow, Count[7:0]=MasterAddress[1:0]+MasterCount[7:0]`

The result of adding `MasterAddress[1:0]` to `MasterCount[7:0]` serves several purposes. First, the overflow field indicates to the requesting processor device that although the size of its request is less than the maximum number of bytes allowed in a transaction, the quadbyte granularity does not allow this to occur and the request should be separated into two separate transactions. Second, the sum produces a count of the number of quadbytes to be transmitted in `Count[7:2]`, which is the granularity of the basic data transport units of the bus. Third, it provides an index in `Count[1:0]` to the last byte to be transported during the last quadbyte of the data packet.

Because the processor supplies the index of the last byte to be transported, the memory device does not need to perform any index arithmetic but instead need only perform a table lookup of the mask data plus a simple logic operation. This reduces the critical path by eliminating the carry chain of the addition. Although the operation is performed by the requesting processor, the processor, unlike the memory device, can typically overlap the addition with other opera-

tions such that the effect is minimized. A significant implementation advantage is achieved which simplifies the receiver of the memory devices by performing the addition at the processor. Typically there are more memory devices than processor devices. It is therefore advantageous to decrease the die size and logic complexity in each of the memory devices in exchange for modestly increasing the complexity of the processor devices to perform this functionality.

The bits Address[1:0] and Count[1:0] are used to generate the masks for the first and last quadbytes of the memory request. The masks are used to determine which bytes within a quadbyte are to be read or written. Masks of varying values are generated only for the first and last quadbytes because all the bytes of the intervening quadbytes will be part of the transaction and the masks therefore have a value of 1111. FIGS. 6a and 6b are tables which respectively illustrate the lookup tables for the mapping of Address[1:0] to Mask[3:0] to generate the mask for the first quadbyte, and the mapping of Count[1:0] to Mask[7:4] to generate the mask for the last quadbyte. A value of one in the mask indicates that the byte is one of the bytes of the memory transaction. Mask[3:0] applies to the first quadbyte at Memory[Address][3:0][8:0]. Mask[7:4] applies to the last quadbyte at Memory[Address+Count][3:0][8:0] ([3:0] identifies the byte of the quadbyte and [8:0] identifies the bit of the byte).

FIGS. 7a-7b illustrate masks generated for byte transfers of various sizes. Referring to FIG. 7a, a single byte transfer is described. A single byte transfer is an illustration of a special case where the first and last quadbyte is the same quadbyte. However, the innovative encoding employed accommodates single quadbyte transfers through simple logic operations which result in simple and space saving logic at the receiver. If Count[7:2] is 00000, the offset-by-one encoding indicates that the transfer is a single quadbyte. When count[7:2] equals 00000, Mask[7:4] and Mask[0:3] fields are logically ANDED together to generate the byte mask for the quadbyte.

FIGS. 7b-7d illustrate the masks generated for, respectively, a two byte transfer, a four byte transfer, and an eight byte transfer. The masks are generated by simple logic bit manipulations which permits simple and fast implementation at the receiver. The arrangement of the bits in the packet are specific to this implementation and lends itself to a space efficient implementation of the logic on the chip. The data sizes correspond to a Count[7:0] value of 00000001, 00000011 and 00000111. For each data size, the four combinations of MasterAddress[1:0] (which is equivalent to the value of Address[1:0]) will be shown, in order of values 00, 01, 10, 11. The use of this encoding and placement of the bits in the packet permit a reasonable compromise between the logic complexity in the processor and the complexity in the memory devices.

Specifically, by placing count bits 6, 4, 2 at bus cycle 4 of the packet and count bits 7, 5, 3 at bus cycle 5 of the packet and respectively on the same signal lines as count bits 6, 4, 2, the amount of wiring needed to interconnect the bits with the logic which processes the count bits is decreased. This saving is reflected in the decrease of the die size. In particular, a carry function is utilized to process the count bits. This is simply and efficiently implemented as bits 2 and 3, 4 and 5, 6 and 7, are aligned, eliminating the need to wire for the carry operation between the bits 2 and 3, 4 and 5, 6 and 7.

While the invention has been described in conjunction with the preferred embodiment, it is evident that numerous alternatives, modifications, variations and uses will be apparent to those skilled in the art in light of the foregoing description.